Slide 1

Significant Advancements in z Systems Resiliency Provided by GDPS

John Thompson
IBM
jgthomps@us.ibm.com

March 5th, 2015
Session 17017

Session QR Code

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides **education**, professional **networking** and industry **influence**.

Disclaimer

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here

Agenda

- GDPS Overview
- GDPS Virtual Appliance
- GDPS/PPRC z/OS Proxy
- GDPS/MTMM (Multi-Target Metro Mirror)

Slide 4

Continuous Availability and Disaster Recovery

First GDPS installation 1998, now more than 770 licenses in 46 countries

- **Central point of control**
  - z Systems and Distributed Servers
  - xDR for z/VM and Linux for z Systems
  - Replication infrastructure
- **Real time monitoring and alert management**
- **Automated Recovery**
  - HyperSwap for Continuous Availability
  - Planned and unplanned outages
- **Configuration Infrastructure Mangement**
  - Single site, 2-site, 3-site, 4-site
- **Automated provisioning**
  - CBU / OOCoD

GDPS
Providing world class protection

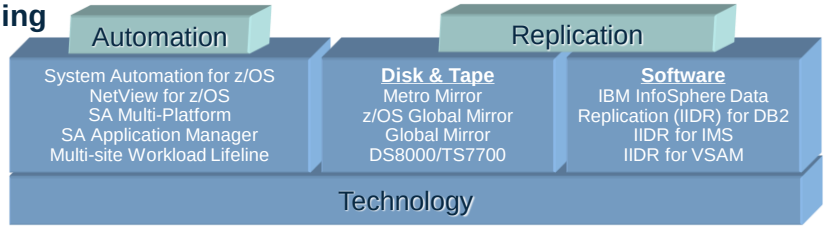

Slide 5

Multiple products meeting various availability objectives

GDPS®/PPRC HM[1]

Near-continuous availability of data within a data center

Single data center

Applications can remain active
Near-continuous access to data in the event of a storage subsystem outage

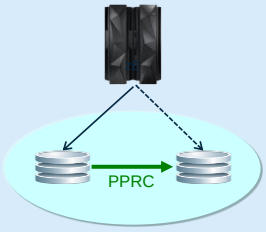

Recovery Point Objective equals 0
and
Recovery Time Objective equals 0

GDPS/PPRC

Near-continuous availability (CA) and disaster recovery (DR) within a metropolitan region

Two data centers

Systems can remain active
Multisite workloads can withstand site and storage failures

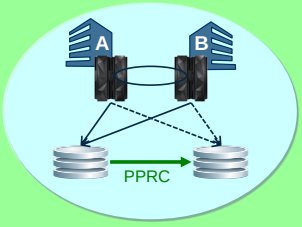

DR RPO equals 0 and RTO is less than 1 hour or
CA RPO equals 0 and RTO minutes

GDPS/GM[2] and GDPS/XRC[3]

Disaster recovery at extended distance

Two data centers

More rapid systems disaster recovery with "seconds" of data loss
Disaster recovery for out-of-region interruptions

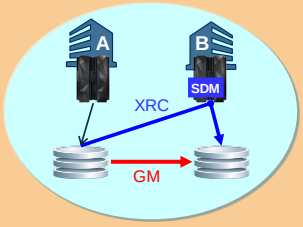

RPO seconds and RTO less than 1 hour

[1]Peer-to-peer remote copy (PPRC), Metro Mirror HyperSwap Manager (HM) [2]Global Mirror (GM) [3]Extended Remote Copy (XRC), zGlobal Mirror

Slide 6

Multiple products meeting various availability objectives (continued)

GDPS/MGM[1] 3-site & GDPS/MzGM[2] 3-site

Near-continuous availability (CA) regionally and disaster recovery (DR) at extended distance

Three data centers

Continuous availability (CA) for site disasters and Disaster recovery (DR) for regional disasters

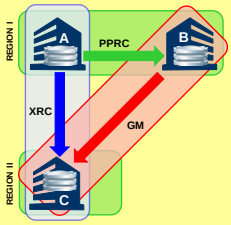

Designed for near-continuous availability within the region and RPO seconds and RTO less then one hour for unplanned region fail-over

GDPS/MGM[1] 4-site & GDPS/MzGM[2] 4-site(SOD)

Near-continuous availability (CA) regionally and disaster recovery (DR) at extended distance from either region

Four data centers

Continuous availability (CA) and Disaster recovery (DR) when running from any site

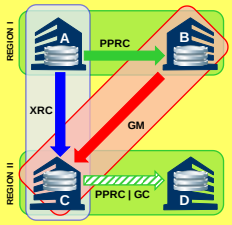

Designed for near-continuous availability within either region and RPO seconds and RTO less then one hour for unplanned region fail-over

GDPS/Active-Active

Near-continuous availability, disaster recovery, and cross-site workload balancing at extended distance

Two or more data centers

Disaster recovery for out-of -region interruptions

All sites active

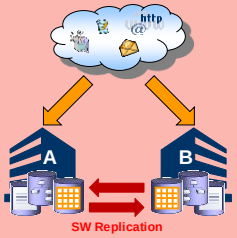

RPO seconds and RTO seconds

[1]Metro Global Mirror (MGM) [2]Metro z/OS Global Mirror (MzGM)

Slide 7

Solutions for z System platforms and for open systems

Multiplatform Resiliency for IBM z Systems (xDR):

- Extends GDPS/PPRC "IBM z/OS® support" to IBM zVM® and associated-guests along with Linux
- Coordinated IBM HyperSwap® – z/OS, z/VM with its guests, and native Linux
- Graceful shutdown and startup of Linux clusters or nodes
- Coordinated takeover – recovery from a Linux cluster failure

Distributed Cluster Manager (DCM):

- Provides near-continuous availability (CA) in both sites
- Extends GDPS to manage front-end distributed servers and clusters (IBM AIX®, HP-UX, Linux, Solaris, VMWare and Microsoft Windows)
- Support for SA AppMan and VCS with Global Cluster Option (GCO)
- Extends monitoring and management to open systems
- Extends data consistency to open systems
- Provides automation scripting for server management

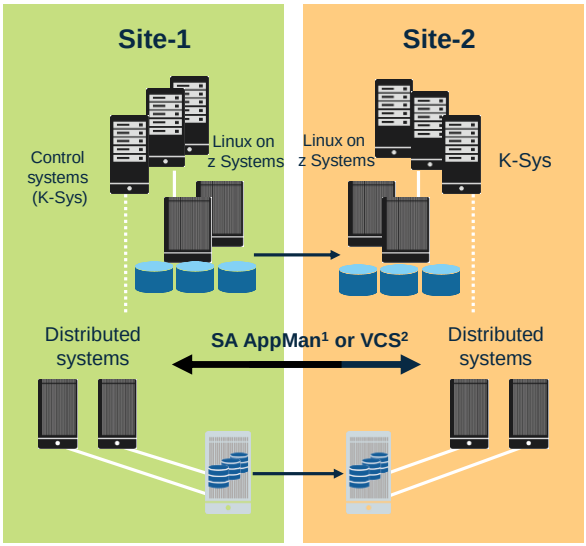

¹IBM Tivoli System Automation Application Manager (SA AppMan), ²Veritas Cluster System (VCS)

Slide 8


SHARE
Educate • Network • Influence
GDPS Virtual Appliance
Session QR Code
CELEBRATING
60
YEARS
OF SHARE
Influencing IT Since 1955
#SHAREorg
SHARE is an independent volunteer-run information technology association
that provides education, professional networking and industry influence.
SHARE
in Seattle 2015
Copyright (c) 2014 by SHARE Inc.
Except where otherwise noted, this work is licensed under
http://creativecommons.org/licenses/by-nc-sa/3.0/

GDPS Virtual Appliance overview

GDPS Virtual Appliance is a fully integrated software solution that provides Continuous Availability & Disaster Recovery functions **for Linux on z Systems with z/VM**

- It is a image comprised of an operating system, the application components, an appliance management layer which makes the image self-containing, and APIs / UIs for customization, administration, and operation tailored to the appliance function
- It extends GDPS/PPRC capabilities into z/VM and Linux on z Systems environments that do not have z/OS in production

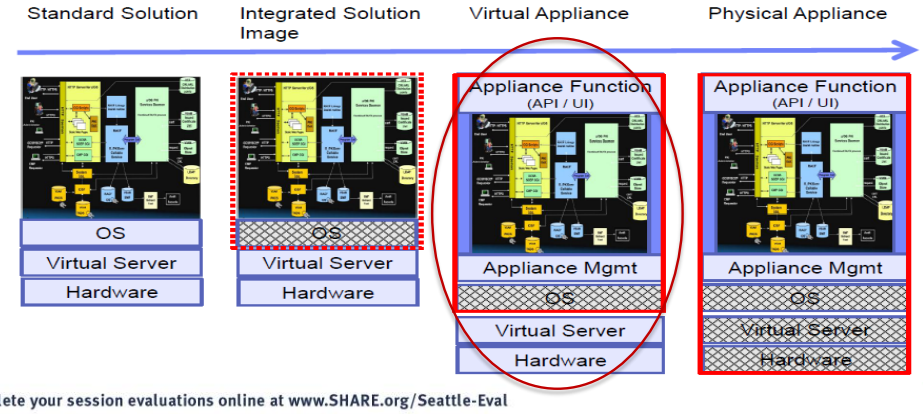

Slide 10

GDPS Virtual Appliance overview (continued)

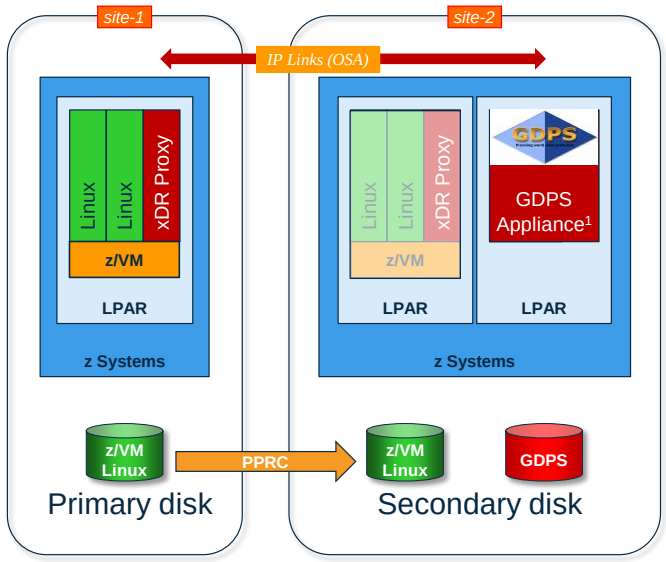

- One dedicated Linux guest is configured as xDR Proxy for GDPS
- GDPS manages PPRC environment and keeps data available & consistent for operating systems and applications
- Disaster detection and assurance of successful and fast recovery via automated processes
- Single point of control from GDPS Appliance. No need for availability of all experts for e.g. storage team, HW team, OS team, application team

[1]pre-configured self-containing image with GDPS K-System

Slide 11

Planned disk outage with GDPS Appliance

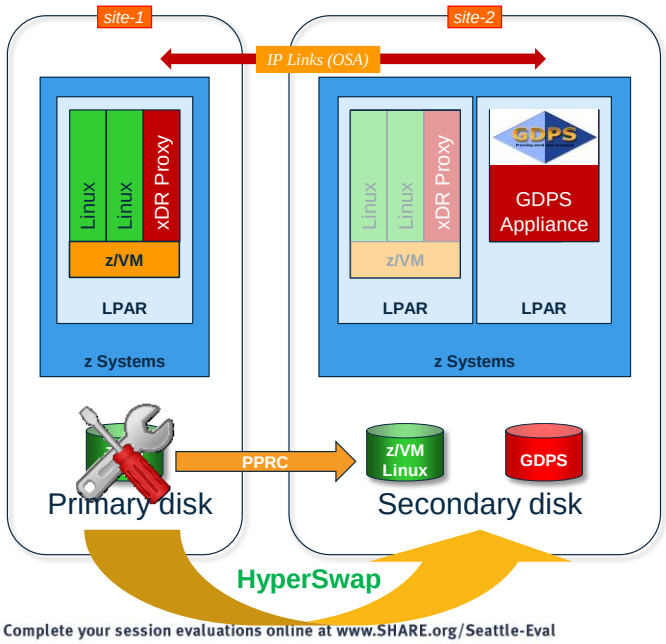

- HyperSwap provides the ability to non-disruptively swap from using the primary volume of a mirrored pair to using what had been the secondary volume
- A planned HyperSwap is invoked manually by operator action using GDPS facilities
- Example: planned HyperSwap initiated in advance of planned disruptive maintenance to a disk subsystem

Unplanned disk outage with GDPS Appliance

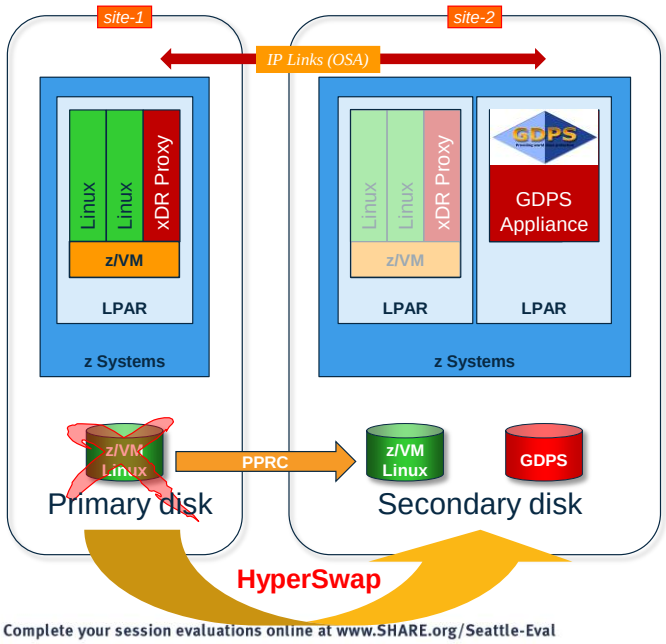

- An unplanned HyperSwap is invoked automatically by GDPS, triggered by events that indicate the failure of a primary disk device
- HyperSwap Triggers
  - I/O errors
  - Boxed devices
  - Control Unit failures
  - Loss of all channel paths
  - "Soft failure" – I/O response time triggers

Disaster Recovery with GDPS Appliance

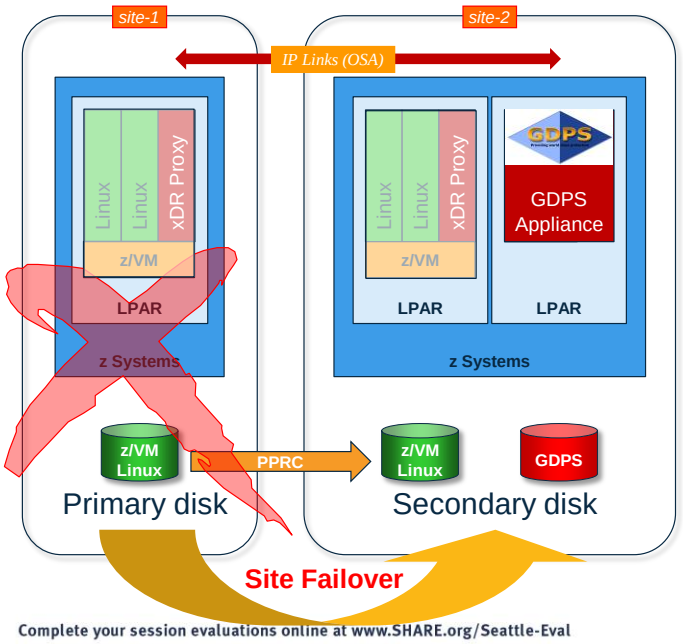

- In case of a site disaster, GDPS Appliance will:
  - immediately issue a freeze for all applicable primary devices to protect the integrity of the secondary data
  - reset Site1 [and Site2] systems and update all the IPL information to point to the new primary devices, and
  - re-IPL all the production systems in LPARs in Site 2
- The GDPS Appliance scripting capability is key to recovering the systems in the shortest possible time following a disaster. All of this is carried out with a single operator instruction

Summary of Requirements

GDPS Appliance hardware requirements:

- 1 LPAR with one single logical General Purpose CP
- 1 GB Memory
- 4 ECKD Disk
- 1 OSA attachment

System prerequisites

- The disks being used by z/VM and Linux to be mirrored must be ECKD disks
- z/VM 5.4 or higher, or z/VM 6.2 or higher
- A supported distribution of Linux on z Systems with latest recommended fixpack
- IBM Tivoli System Automation for Multiplatforms with latest recommended fixpack
  - Separately priced xDR for Linux feature is required

Supported Linux distributions	TSAMP Release Level
SUSE SLES 11	SA MP 4.1 or higher
Red Hat 6	SA MP 4.1 or higher

Slide 15


SHARE
Educate • Network • Influence
GDPS/PPRC z/OS Proxy
Session QR Code
CELEBRATING
60
YEARS
OF SHARE
Influencing IT Since 1955
#SHAREorg
SHARE is an independent volunteer-run information technology association
that provides education, professional networking and industry influence.
SHARE
in Seattle 2015
Copyright (c) 2014 by SHARE Inc.
Except where otherwise noted, this work is licensed under
http://creativecommons.org/licenses/by-nc-sa/3.0/

Slide 16

Multiple Sysplexes and monoplex systems

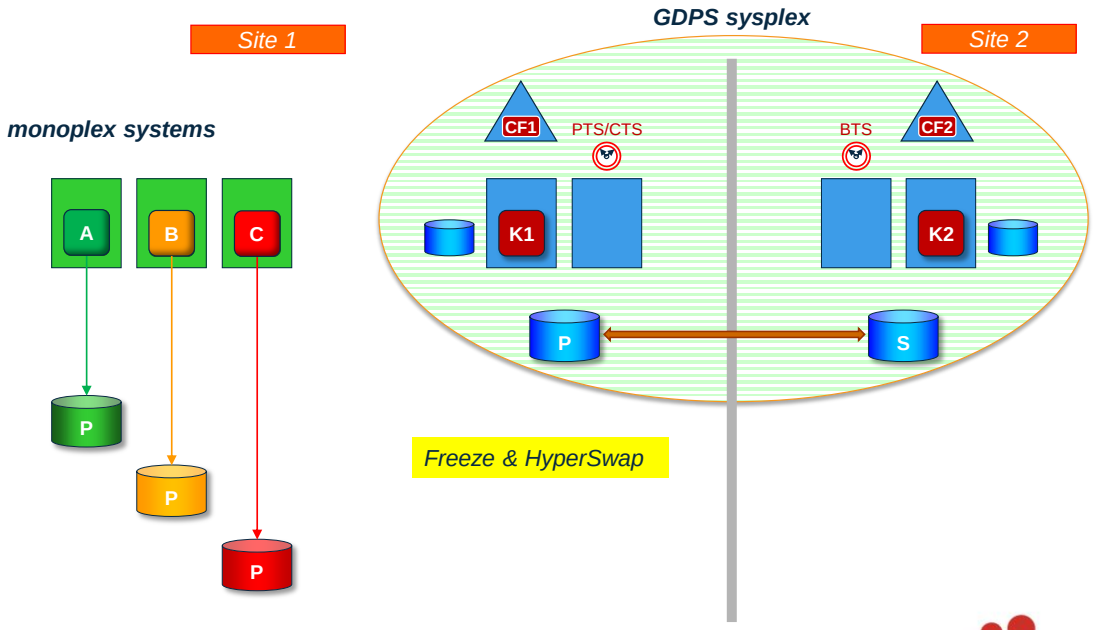

Slide 17

GDPS/PPRC z/OS Proxy

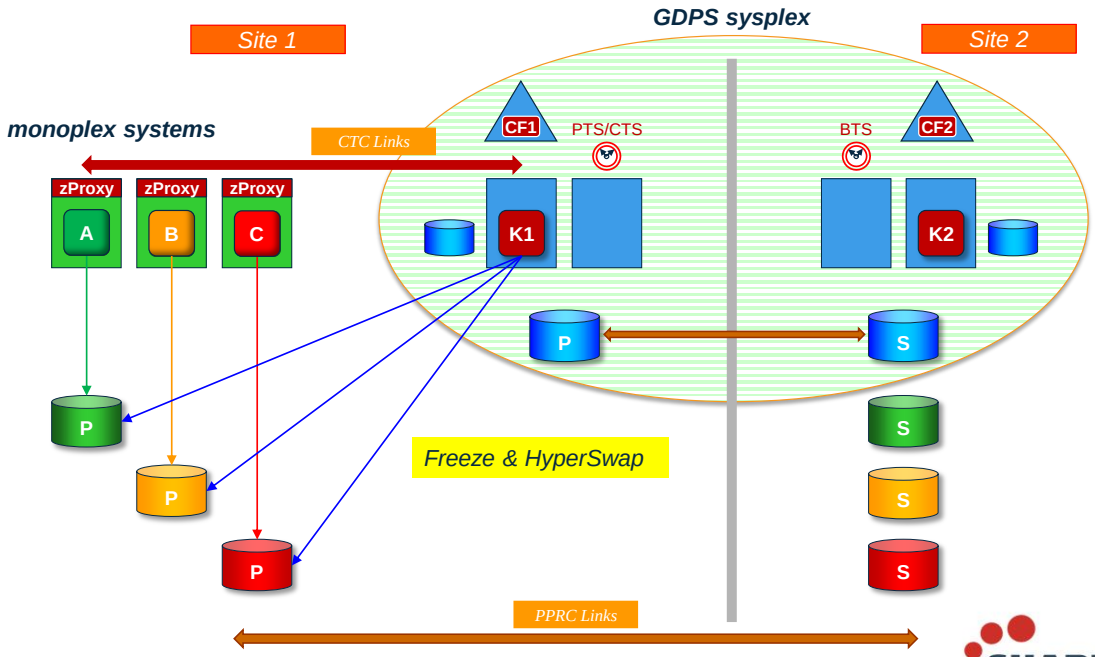

z/OS Proxy Basics

- Separately priced feature of GDPS/PPRC (GDPS/PPRC only)
- Max 16 foreign systems monitored simultaneously – z/VM, Linux systems not included
- Communication to z/OS Proxy managed system(s) – CTC based (Channel To Channel)
  - xDR Communication to z/VM and Linux on z Systems – IP based
- No disk sharing between multiple sysplexes or between sysplexes and stand-alone systems

Key benefits…

- Satisfies requirements for K System consolidation
- Minimizes number of configurations/environments to manage and maintain

Slide 19


SHARE
Educate • Network • Influence
GDPS/MTMM (Multi-Target Metro Mirror)
Session QR Code
CELEBRATING
60
YEARS
OF SHARE
Influencing IT Since 1955
#SHAREorg
SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.
SHARE in Seattle 2015
Copyright (c) 2014 by SHARE Inc.
Except where otherwise noted, this work is licensed under http://creativecommons.org/licenses/by-nc-sa/3.0/

Slide 20

MultiTarget Metro Mirror (MTMM)

- Data is transferred to both targets in parallel
- Pairs operate independent of each other
- Pairs may be established, suspended or removed separately
- A replication problem on one leg does not impact the other leg
- HyperSwap is possible on either leg

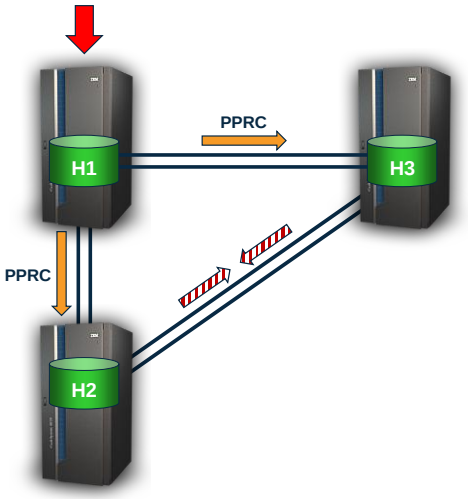

Provides all the benefits of PPRC synchronous mirroring plus has the additional protection of a second synchronous leg

- A typical configuration has the H2 more local than H3.
  - H2 provides for local recovery in case of a failure of H1 (Continuous Availability - CA)
  - H3 provides for a more remote recovery in the case of a more widespread failure that affects both H1 and H2 (Disaster Recovery – DR)

Recovery from a primary disk failure

- Failure at H1
- HyperSwap to H2
  - Failover H2→H1
  - Move I/O to H2
- Resume H2→H3 with Incremental Resync (IR)
  - Failover H3→H1
  - Failback H2→H3
  - Mirroring is quickly resumed
- H1 recovered
  - Failback H2→H1

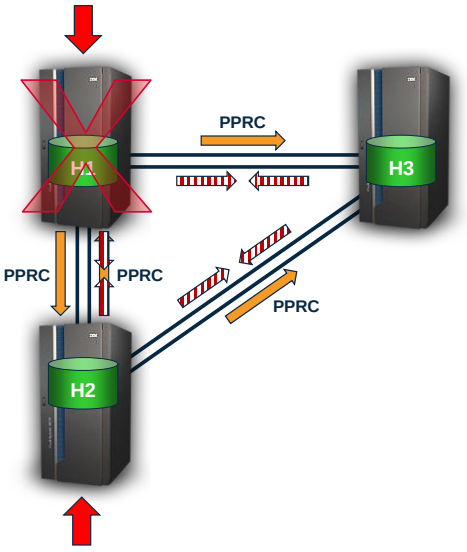

Each replication leg is tracked and managed independently

Return to normal

- HyperSwap to H1
  - Failover H1→H2
  - Move I/O to H1
- Resume H1→H2 and H1→H3 pairs
  - Failover H3→H2
  - Failback H1→H2
  - Failback H1→H3

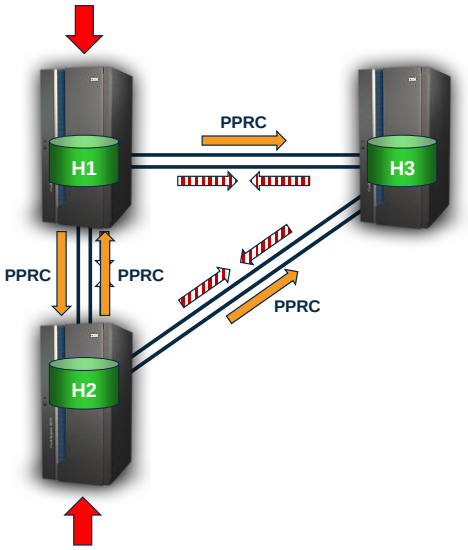

Maintaining two PPRC secondary targets (two PPRC legs) from a single primary device

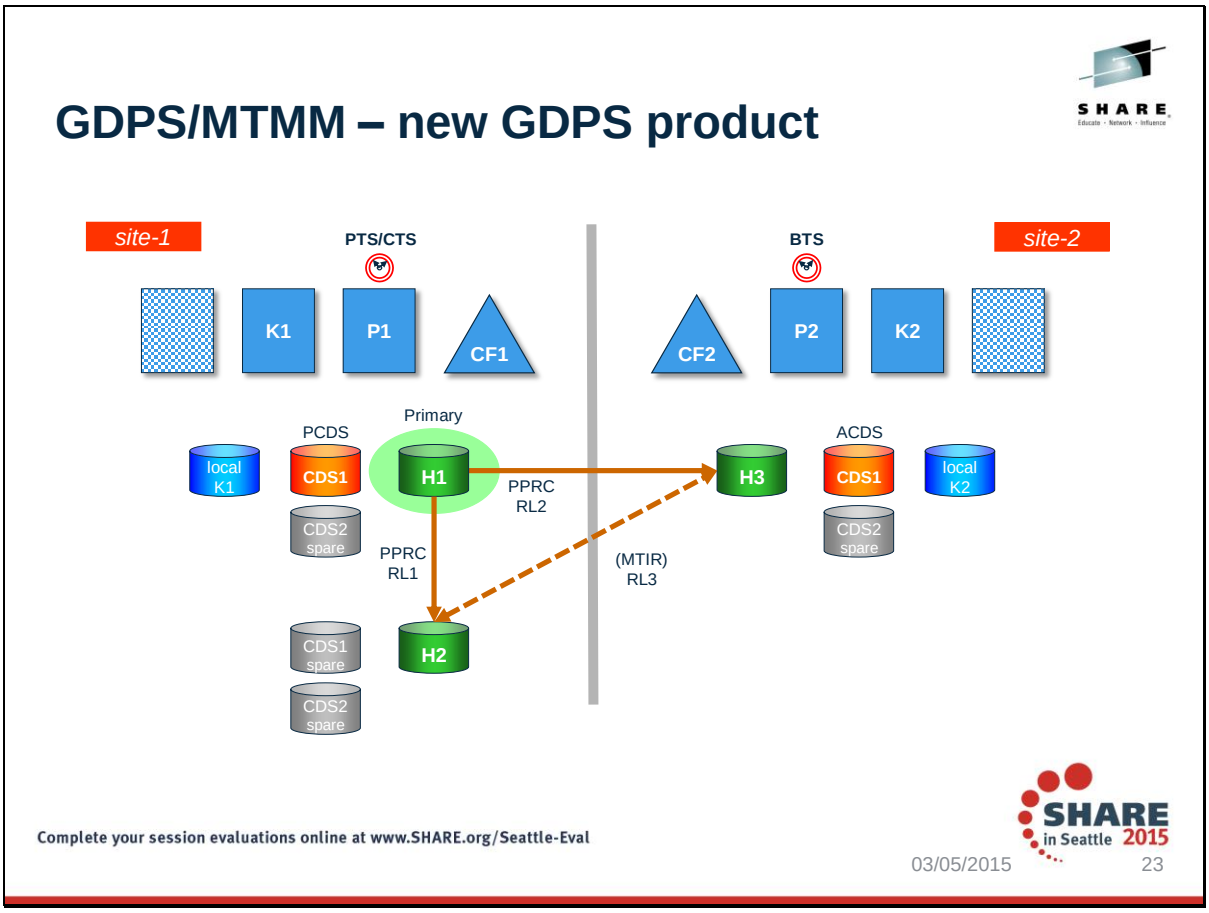

- New terminology:
  - "Site 1 disk" now ambiguous
  - Three "disk locations:" H1, H2, and H3
  - Three "replication legs:" RL1, RL2, and RL3. The replication leg between the two current secondary disks is known as the "MultiTarget Incremental Resync" or "MTIR" leg.
- Two sets of non-PPRCed CDSs in each of the disk locations. The general concepts of Normal, Site1-only, and Site2-only configurations in GDPS/PPRC also exist in GDPS/MTMM, however the concepts are not identical across the two products. For GDPS/MTMM:
  - **Normal** - CDS1 of the current primary disk storage location will be used as the Primary CDS. The Alternate CDS will be the CDS1 defined in the secondary disk storage location that provides failure isolation at a site level. This means that H3 would contain CDS2 when the primary disks are running in Site1 (H1 or H2). When the primary disks are in H3,

the CDS1 defined in the secondary storage location of the current preferred swap leg would be used as the Alternate.

- **Site1** Uses CDS1 of current primary disk storage location as the Primary CDS if the primary disk storage is in H1 or H2. CDS1 of the other storage location would be used as the Alternate CDS. If the primary disk is in H3, then the CDS1 defined in the secondary storage location of the current preferred swap leg would be used as the Primary CDS, and the CDS1 defined in the other storage Site1 location would be used as the Alternate.

- **Site2** CDS1 and CDS2 of the H3 storage location would be used as the Primary and Alternate CDSs respectively.

- Although not shown, there are a number of cross-site connections between the two sites over dark, dedicated fiber, using DWDMs. There is connectivity for various host to disk FICON channel connections, PPRC links, sysplex connections for XCF signaling amongst the systems and Server Time Protocol connectivity. Additionally all servers in both sites are connected to the same HMC/SE LAN requiring the HMC/SE LAN to be cross-connected either over the DWDM fiber infrastructure or bridged over the WAN.

Key Features of new Product

- Provides the ability to HyperSwap on either or both of the PPRC legs
  - Planned HyperSwap targets the leg specified, other leg suspends
  - User policy to specify which leg is considered the ‘preferred’ swap leg for unplanned HyperSwap
  - Unplanned HyperSwap will occur on preferred leg if possible else, other leg attempted
- Provides ability to FREEZE impacted leg only on PPRCFAILURE failures
  - Policy options provided to specify actions to take (STOP or COND) ***when the last active leg suspends*** (action will be GO when the other leg is still duplex)
- Provides monitoring and automation to ensure reliable and rapid site switching via automated processes

Key Features of new Product (continued)

- Provides CDS management capabilities that fully support 3 storage locations
- Majority of system, sysplex and HW management and monitoring functions inherited from GDPS/PPRC
- z/VM and Linux for z Systems guests running under z/VM are supported (xDR)

Summary of Requirements and Prerequisites

Storage environment

- All disk subsystems in the configuration must be DS8870 with MC level 7.4
- Only CKD disk is supported
- FCP links mandatory

Systems and software

- z/OS 1.13 or 2.1 (or higher) plus related DFSMS, IOS, zVM PTFs
- NetView for z/OS V5.3 or V6.1 (or higher) or IBM Service Management Suite for z/OS V1.1 (or higher)
- Tivoli System Automation for z/OS V3.4 (or higher) or IBM Service Management Suite for z/OS V1.1 (or higher)

Summary of Requirements and Prerequisites (continued)

xDR Environments

- z/VM 5.4, or z/VM 6.2 or higher
- A supported distribution of Linux on z Systems with latest recommended fixpack
- IBM Tivoli System Automation for Multiplatforms with latest recommended fixpack
  - Separately priced xDR for Linux feature is required

IBM GDPS solutions – Value for your organization

Experience	Commitment	Value	Vision

Customer acceptance

- More than 770 IBM GDPS® licenses installed in 46 countries worldwide
- Proven technology, automated and repeatable result
- Complete implementation by experienced consultants

Open industry standards

- GDPS supports industry-accepted, open replication architectures (PPRC[1], XRC[2], Global Mirror and Fibre Channel)
- Architectures licensed by all enterprise storage vendors
- GDPS qualification program (IBM and Hitachi)

Investment protection

- Designed to be easily upgradeable
- Common code base for each product

Product maturity

- Generally available since 1998
- Suite of products
- Enterprise-to-enterprise capability
- Many years of IBM z Systems® production experience
- CA and DR[3] best of breed
- Continually enhanced

Customer focus

- GDPS Design Council
- Synergy with IBM development labs
- Incorporates several IBM patents
- New release planned every year

IBM support

- Fully supported via standard IBM support structure
- Fixes through normal IBM z Systems® channels

Additional Information

- **Web sites:**

GDPS	**www.ibm.com/systems/z/gdps**
Parallel Sysplex	**www.ibm.com/systems/z/pso**
Bus Resiliency z	**www.ibm.com/systems/z/resiliency**
Bus Resiliency	**www.ibm.com/systems/business_resiliency**

- z Systems **www.ibm.com/systems/z/hardware**
  Storage **www.ibm.com/systems/storage**
  Redbooks®GDPS Family: An Introduction to Concepts and Capabilities
  www.redbooks.ibm.com/abstracts/sg246374.html?Open

- **GDPS Web Site White Papers and Presentations**
  - GDPS: The Ultimate e-business Availability Solution
  - IBM Implementation Services for GDPS/Global Mirror
  - GDPS Business Continuity Solutions
  - Consistency Groups in a Nutshell
  - DS8000™ Data Replication
  - GDPS Solutions

- **e-mail:** gdps@us.ibm.com